

REPORT DOCUMENTATION PAGE			Form Approved OMB NO. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>				
1. REPORT DATE (DD-MM-YYYY) 10-06-2015		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 1-Mar-2011 - 28-Feb-2015
4. TITLE AND SUBTITLE Final Report: Dynamic Routing and Coordination in Multi-Agent Networks			5a. CONTRACT NUMBER W911NF-11-1-0092	
			5b. GRANT NUMBER	
			5c. PROGRAM ELEMENT NUMBER 611102	
6. AUTHORS Francesco Bullo			5d. PROJECT NUMBER	
			5e. TASK NUMBER	
			5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES University of California - Santa Barbara 3227 Cheadle Hall 3rd floor, MC 2050 Santa Barbara, CA 93106 -2050			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS (ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211			10. SPONSOR/MONITOR'S ACRONYM(S) ARO	
			11. SPONSOR/MONITOR'S REPORT NUMBER(S) 58305-NS.29	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for Public Release; Distribution Unlimited				
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.				
14. ABSTRACT Supported by this project, we designed innovative routing, planning and coordination strategies for robotic networks and studied their application to Army and DoD scenarios. The key technological challenge is the decision of who does what, when or, equivalently, how are tasks partitioned among robots, in what order are they to be performed, and along which deterministic routes or according to which stochastic rules do individual robots move. The fundamental novelties and our recent breakthroughs supported by this project are manifold: (1) the application				
15. SUBJECT TERMS dynamic vehicle routing, robotic coordination, controllability on complex networks, robotic surveillance				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU		
				19a. NAME OF RESPONSIBLE PERSON Francesco Bullo
				19b. TELEPHONE NUMBER 805-893-5169

Report Title

Final Report: Dynamic Routing and Coordination in Multi-Agent Networks

ABSTRACT

Supported by this project, we designed innovative routing, planning and coordination strategies for robotic networks and studied their application to Army and DoD scenarios. The key technological challenge is the decision of who does what, when or, equivalently, how are tasks partitioned among robots, in what order are they to be performed, and along which deterministic routes or according to which stochastic rules do individual robots move.

The fundamental novelties and our recent breakthroughs supported by this project are manifold: (1) the application of queueing theory and combinatorial techniques to network of autonomous robots leads to a wide range of new relevant problems, (2) novel coordination schemes promise to successfully achieve various optimization objectives relying only upon asynchronous and asymmetric communication, (3) the increasingly weaker assumptions imposed on routing and coordination algorithms are rendering them practical and widely applicable.

This project addressed multi-dimensional problems of relevance in Engineering and Computer Science by unifying fundamental concepts from multiple domains (robotics, autonomy, combinatorics, and network science). Our work aimed to bridge multiple scientific disciplines, including control theory and theoretical computer science and their applications to multi-agent systems, robotics and sensor networks.

Technology Transfer

Interactions with United Technology Research Center on the topic of Autonomy and Cognition: Dr Luca Bertuccelli and Amit Surana visited UCSB; UCSB students Vaibhav Srivastava and Jeffrey Peters visited UTRC. Collaborative effort on Supervisory Controller for Optimal Role Allocation for Cueing of Human Operators has been funded through the ICB 6.2 program.

Participation to the ARO Network Sciences Strategy Planning Workshop, on January 16-17, 2014 at the Adelphi Laboratory Center in Maryland.

Participation to the ARO Workshop on Emerging Cyber War-Fighting Technologies, on March 10-11, 2015 in College Park, Maryland.

Research Overview and detailed Presentation during UCSB visit by Dr. Purush Iyer on December 3, 2014.

Enter List of papers submitted or published that acknowledge ARO support from the start of the project to the date of this printing. List the papers, including journal references, in the following categories:

(a) Papers published in peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
03/07/2012 1.00	Francesco Bullo, Emilio Frazzoli, Marco Pavone, Ketan Savla, Stephen L. Smith. Dynamic Vehicle Routing for Robotic Systems, Proceedings of the IEEE, (09 2011): 0. doi:
03/07/2012 2.00	Mathias Burger, Giuseppe Notarstefano, Francesco Bullo, Frank Allgower. A Distributed Simplex Algorithm for Degenerate Linear Programs and Multi-Agent Assignment, Automatica, (07 2011): 0. doi:
08/11/2014 14.00	Fabio Pasqualetti, Domenica Borra, Francesco Bullo. Consensus networks over finite fields, Automatica, (02 2014): 0. doi: 10.1016/j.automatica.2013.11.011
08/11/2014 15.00	Fabio Pasqualetti, Sandro Zampieri, Francesco Bullo. Controllability Metrics, Limitations and Algorithms for Complex Networks, IEEE Transactions on Control of Network Systems, (03 2014): 0. doi: 10.1109/TCNS.2014.2310254
08/11/2014 16.00	Filippo Zanella, Jeffrey Russel Peters, Markus Spindler, Fabio Pasqualetti, Ruggero Carli, Francesco Bullo. Camera Network Coordination for Intruder Detection, IEEE Transactions on Control Systems Technology, (09 2014): 0. doi: 10.1109/TCST.2013.2290708
08/11/2014 17.00	Rushabh Patel, Paolo Frasca, Francesco Bullo. Centroidal Area-Constrained Partitioning for Robotic Networks, Journal of Dynamic Systems, Measurement, and Control, (03 2014): 0. doi: 10.1115/1.4026344
08/11/2014 18.00	Vaibhav Srivastava, Francesco Bullo. Knapsack problems with sigmoid utilities: Approximation algorithms via hybrid optimization, European Journal of Operational Research, (07 2014): 0. doi: 10.1016/j.ejor.2013.12.035
08/11/2014 19.00	Sara Susca, Pushkarini Agharkar, Sonia Martínez, Francesco Bullo. Synchronization of Beads on a Ring by Feedback Control, SIAM Journal on Control and Optimization, (03 2014): 0. doi: 10.1137/120903208
08/11/2014 20.00	V. Srivastava, F. Pasqualetti, F. Bullo. Stochastic surveillance strategies for spatial quickest detection, The International Journal of Robotics Research, (09 2013): 0. doi: 10.1177/0278364913490322
08/11/2014 21.00	Mauro Franceschelli, Daniele Rosa, Carla Seatzu, Francesco Bullo. Gossip algorithms for heterogeneous multi-vehicle routing problems, Nonlinear Analysis: Hybrid Systems, (11 2013): 0. doi: 10.1016/j.nahs.2013.03.001
08/11/2014 22.00	Domenica Borra, Fabio Pasqualetti, Francesco Bullo. Continuous Graph Partitioning for Camera Network Surveillance, Automatica, (04 2014): 0. doi:
08/21/2012 3.00	Fabio Pasqualetti, Joseph W. Durham, Francesco Bullo. Cooperative Patrolling via Weighted Tours: Performance Analysis and Distributed Algorithms, IEEE Transactions on Robotics, (11 2011): 0. doi:

TOTAL: 12

Number of Papers published in peer-reviewed journals:

(b) Papers published in non-peer-reviewed journals (N/A for none)

<u>Received</u>	<u>Paper</u>
08/21/2012 4.00	Luca Carlone, Vaibhav Srivastava, Francesco Bullo, Giuseppe Calafiore. DISTRIBUTED RANDOM CONVEX PROGRAMMING VIA CONSTRAINTS CONSENSUS, SIAM J of Control and Optimization, (07 2012): 0. doi:
TOTAL:	1

Number of Papers published in non peer-reviewed journals:

(c) Presentations

Number of Presentations: 0.00

Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
TOTAL:	

Number of Non Peer-Reviewed Conference Proceeding publications (other than abstracts):

Peer-Reviewed Conference Proceeding publications (other than abstracts):

<u>Received</u>	<u>Paper</u>
TOTAL:	

(d) Manuscripts

<u>Received</u>	<u>Paper</u>
06/09/2015 26.00	R. Patel, A. Carron, F. Bullo. The hitting time of multiple random walks with application to robotic surveillance, SIAM Journal on Matrix Analysis and Applications (03 2015)
06/09/2015 28.00	S. Bopardikar, F. Bullo, P. Agharkar. Vehicle routing algorithms for radially escaping targets, SIAM Journal on Control and Optimization (12 2014)
06/09/2015 27.00	P. Agharkar, F. Bullo. Quickest detection over robotic roadmaps, IEEE Transactions on Robotics (04 2015)
08/11/2014 23.00	Rushabh Patel, Pushkarini Agharkar, Francesco Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant, IEEE Transactions on Automatic Control (04 2014)
08/11/2014 24.00	Rushabh Patel, Paolo Frasca, Joseph W. Durham, Ruggero Carli, Francesco Bullo. Dynamic Partitioning and Coverage Control with Asynchronous One-to-Base-Station Communication, IEEE Transactions on Control of Network Systems (01 2014)
08/21/2012 5.00	Vaibhav Srivastava, Fabio Pasqualetti, Francesco Bullo. Stochastic Surveillance Strategies for Spatial Quickest Detection, International Journal of Robotics Research (04 2012)
08/21/2012 6.00	Domenica Borra, Fabio Pasqualetti, Francesco Bullo. Continuous Graph Partitioning for Camera Network Surveillance, Automatica (07 2012)
08/28/2013 7.00	Shaunak D. Bopardikar, Stephen L. Smith, Francesco Bullo. On Dynamic Vehicle Routing with Time Constraints, IEEE Transactions on Robotics (07 2013)
08/28/2013 8.00	Rush Patel, Paolo Frasca, Francesco Bullo. Centroidal Area-Constrained Partitioning for Robotic Networks, ASME Journal of Dynamic Systems, Measurement, and Control (03 2013)
08/28/2013 9.00	Fabio Pasqualetti, Domenica Borra, Francesco Bullo. Consensus Networks over Finite Fields, Automatica (01 2013)
08/28/2013 10.00	Mauro Franceschelli, Daniele Rosa, Carla Seatzu, Francesco Bullo. Gossip Algorithms for Heterogeneous Multi-Vehicle Routing Problems, Nonlinear Analysis: Hybrid Systems (10 2012)
08/28/2013 11.00	Vaibhav Srivastava, Francesco Bullo. Knapsack Problems with Sigmoid Utilities: Approximation Algorithms via Hybrid Optimization, European Journal of Operational Research (10 2012)
08/28/2013 12.00	Fabio Pasqualetti, Filippo Zanella, Jeffrey R. Peters, Markus Spindler, Ruggero Carli, Francesco Bullo. Camera Network Coordination for Intruder Detection, IEEE Transactions on Control Systems Technology (01 2013)

08/28/2013 13.00 Sara Susca, Pushkarini Agharkar, Sonia Martinez, Francesco Bullo. SYNCHRONIZATION OF BEADS
ON A RING BY FEEDBACK CONTROL,
SIAM Journal on Control and Optimization (12 2012)

TOTAL: 14

Number of Manuscripts:

Books

Received Book

TOTAL:

Received Book Chapter

TOTAL:

Patents Submitted

Patents Awarded

Awards

IFAC Automatica Best Paper Award, 2014

SIAG/CST Best Paper Prize, SIAM Journal on Control and Optimization, 2013

O. Hugo Schuck Best Paper Award, American Automatic Control Council, 2011

Invited Survey, Automatica, 50th anniversary volume, 2014

Article in Inaugural Issue, IEEE Transactions on Control of Network Systems, 2014

(acceptance rate: 11 articles out of 120 submissions)

Distinguished Lecture Series in Cyberphysical Systems, University of New Mexico, Jan 2014

Plenary Speaker:

16th Latin American Congress of Automatic Control, Cancún, México, Oct 2014

11th Int. Symp. on Distributed Autonomous Robotic Systems (DARS), Baltimore, MD, USA, Nov 2012 5th Georgia Tech

Decision & Control Student Symposium, Atlanta, GA, USA, Apr 2012

11th SIAM Conference on Control & Its Applications, Baltimore, MD, USA, Jul 2011

SemiPlenary or Keynote Speaker:

20th Int. Symp. on Mathematical Theory of Networks and Systems (MTNS), Melbourne, Australia, Jul 2012

Outstanding Graduate Mentor Award, UCSB Academic Senate, 2015

Graduate Students

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	Discipline
Jeff Peters	0.10	
Vaibhav Srivastava	0.25	
Pushkarini Agharkar	0.50	
Rush Patel	0.10	
John Simpson	0.10	
Sepher Seifi	0.10	
FTE Equivalent:	1.15	
Total Number:	6	

Names of Post Doctorates

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
Fabio Pasqualetti	0.15
FTE Equivalent:	0.15
Total Number:	1

Names of Faculty Supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>	National Academy Member
Francesco Bullo	0.10	
FTE Equivalent:	0.10	
Total Number:	1	

Names of Under Graduate students supported

<u>NAME</u>	<u>PERCENT SUPPORTED</u>
FTE Equivalent:	
Total Number:	

Student Metrics

This section only applies to graduating undergraduates supported by this agreement in this reporting period

The number of undergraduates funded by this agreement who graduated during this period: 0.00

The number of undergraduates funded by this agreement who graduated during this period with a degree in science, mathematics, engineering, or technology fields:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and will continue to pursue a graduate or Ph.D. degree in science, mathematics, engineering, or technology fields:..... 0.00

Number of graduating undergraduates who achieved a 3.5 GPA to 4.0 (4.0 max scale):..... 0.00

Number of graduating undergraduates funded by a DoD funded Center of Excellence grant for Education, Research and Engineering:..... 0.00

The number of undergraduates funded by your agreement who graduated during this period and intend to work for the Department of Defense 0.00

The number of undergraduates funded by your agreement who graduated during this period and will receive scholarships or fellowships for further studies in science, mathematics, engineering or technology fields:..... 0.00

Names of Personnel receiving masters degrees

NAME

Total Number:

Names of personnel receiving PHDs

NAME

Rush Patel

Vaibhav Srivastava

Total Number: 2

Names of other research staff

NAME

PERCENT SUPPORTED

FTE Equivalent:

Total Number:

Sub Contractors (DD882)

Inventions (DD882)

Scientific Progress

Scientific Progress and Accomplishments
Agreement Number: ARO W911NF-11-1-0092
"Dynamic Routing and Coordination in Multi-Agent Networks"
Final Report for comprehensive period: March 2011 - February 2015

Francesco Bullo
Department of Mechanical Engineering
University of California, Santa Barbara

June 8, 2015

Abstract

We propose the design of innovative routing, planning and coordination strategies for robotic networks and their application to Army and DoD scenarios. The key technological challenge is the decision of who does what, when or, equivalently, how are tasks partitioned among robots, in what order are they to be performed, and along which routes do individual robots move.

The fundamental novelties and our recent breakthroughs motivating this proposal are manifold: (1) the application of queueing theory and combinatorial techniques to network of autonomous robots leads to a wide range of new relevant problems, (2) novel coordination schemes promise to successfully achieve various optimization objectives relying only upon asynchronous and asymmetric communication, (3) the increasingly weaker assumptions imposed on routing and coordination algorithms are rendering them practical and widely applicable.

This proposal addresses multi-dimensional problems of relevance in Engineering and Computer Science by unifying fundamental concepts from multiple domains (robotics, autonomy, combinatorics, and network science). The proposal aims to bridge multiple scientific disciplines, including control theory and theoretical computer science and their applications to multi-agent systems, robotics and sensor networks.

Throughout the duration of this grant, we made solid contributions in numerous key areas:

1. stochastic surveillance for quickest detection,
2. dynamic vehicle routing,
3. coverage control and equitable partitioning,
4. knapsack optimization with application to mixed teams,
5. controllability for complex networks,
6. cooperative patrolling for robotic and camera networks,
7. distributed optimization and task assignment

In the following, we review our scientific progress and accomplishments in each of these areas.

1 Multi-robot stochastic surveillance for quickest detection

In this first work (see (J1) below), we design persistent surveillance strategies for the quickest detection of anomalies taking place in an environment of interest. From a set of predefined regions in the environment, a team of autonomous vehicles collects noisy observations, which a control center processes. The overall objective is to minimize detection delay while maintaining the false alarm rate below a desired threshold. We present joint (i) anomaly detection algorithms for the control center and (ii) vehicle routing policies. For the control center, we propose parallel cumulative sum (CUSUM) algorithms (one for each region) to detect anomalies from noisy observations. For the vehicles, we propose a stochastic routing policy, in which the regions to be visited are chosen according to a probability vector. We study stationary routing policy (the probability vector is constant) as well as adaptive routing policies (the probability vector varies in time as a function of the likelihood of regional anomalies). In the context of stationary policies, we design a performance metric and minimize it to design an efficient stationary routing policy. Our adaptive policy improves upon the stationary counterpart by adaptively increasing the selection probability of regions with high likelihood of anomaly. Finally, we show the effectiveness of the proposed algorithms through numerical simulations and a persistent surveillance experiment.

In our more-recent work (see (J2) below), we provide analysis and optimization results for the *mean first passage time*, also known as the *Kemeny constant*, of a Markov chain. First, we generalize the notion of Kemeny constant to environments with heterogeneous travel and service times, denote this generalization as the *weighted Kemeny constant*, and we characterize its properties. Second, for reversible Markov chains, we show that the minimization of the Kemeny constant and its weighted counterpart can be formulated as convex optimization problems and, moreover, as semidefinite programs. Third, we apply these results to the design of stochastic surveillance strategies for quickest detection of anomalies in network environments. We numerically illustrate the proposed design: compared with other well-known Markov chains, the performance of our Kemeny-based strategies are always better and in many cases substantially so. These results leave open numerous directions for further research. First, we designed surveillance policy only for single agent systems and it would be of practical interest to consider the case where there are multiple agents. Second, it would be useful to understand bounds on the design of the mean first passage time for general graph topologies. Finally, we treat only the optimization of the transition matrix of the graph. It would be of interest to study how we can optimize the weight matrix W in conjunction with the transition matrix. This can have the interpretation of optimizing the "capacity" or "resistance" of the graph, a topic in optimization which is of independent interest.

Finally, in (J3), we provide generalized notions and analysis methods for the *hitting time* of random walks on graphs. The hitting time, also known as the Kemeny constant or the mean first passage time, of a random walk is widely studied, however, only limited work is available for the multiple random walker scenario. In this work we provide a novel method for calculating the hitting time for a single random walker as well as the first analytic expression for calculating the hitting time for multiple random walkers, which we denote as the *group hitting time*. We also provide closed form solution for calculating the hitting time between specified nodes for both the single and multiple random walker cases. Our results allow for the multiple random walks to be different and, moreover, for the random walks to operate on different subgraphs. Finally, using sequential quadratic programming, we show the combination of transition matrices that generate the minimal group hitting time for various graph topologies are often different.

Bibliography

- (J1) V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, 32(12):1438–1458, 2013
- (J2) R. Patel, P. Agharkar, and F. Bullo. Robotic surveillance and Markov chains with minimal first passage time. *IEEE Transactions on Automatic Control*, May 2014. To appear
- (J3) R. Patel, A. Carron, and F. Bullo. The hitting time of multiple random walks with application to robotic surveillance. *SIAM Journal on Matrix Analysis and Applications*, March 2015. Submitted
- (J4) P. Agharkar and F. Bullo. Quickest detection over robotic roadmaps. *IEEE Transactions on Robotics*, April 2015. Submitted
- (PhD) R. Patel. *Robotic Surveillance and Deployment Strategies*. PhD thesis, Mechanical Engineering Department, University of California at Santa Barbara, April 2015

2 Dynamic vehicle routing

In this thrust, we completed the survey submission (J1) below. This paper surveys recent concepts and algorithms for dynamic vehicle routing (DVR), that is, for the automatic planning of optimal multi-vehicle routes to perform tasks that are generated over time by an exogenous process. We consider a rich variety of scenarios relevant for robotic applications. We begin by reviewing the basic DVR problem: demands for service arrive at random locations at random times and a vehicle travels to provide on-site service while minimizing the expected wait time of the demands. Next, we treat different multi-vehicle scenarios based on different models for demands (e.g., demands with different priority levels and impatient demands), vehicles (e.g., motion constraints, communication and sensing capabilities), and tasks. The performance criterion used in these scenarios is either the expected wait time of the demands or the fraction of demands serviced successfully. In each specific DVR scenario, we adopt a rigorous technical approach that relies upon methods from queueing theory, combinatorial optimization and stochastic geometry. First, we establish fundamental limits on the achievable performance, including limits on stability and quality of service. Second, we design algorithms, and provide provable guarantees on their performance with respect to the fundamental limits.

In the article (J2), we considered the problem of dynamic vehicle routing under exact time constraints on servicing demands. Demands for service are generated in an environment as follows: uniformly randomly in space and Poisson in time. Every demand needs to be serviced exactly after a fixed, finite interval of time after it is generated. We design routing policies for a service vehicle to maximize the fraction of demands serviced at steady-state. The main contributions are as follows. First, we demonstrate that this problem is described by an appropriate directed acyclic graph structure which leads to a computationally-efficient routing algorithm based on a longest-path computation. Second, we provide two analytic lower bounds on the service fraction of the longest path policy. The first bound is relative to an optimal, non-causal version of the policy, i.e., a policy based on knowledge of all future demand requests. The second bound is an explicit function of demand generation rate, and therefore, useful as a design tool. We also present numerical results to support the analytic bounds and to shed light on parameter regimes where the analytic bounds are not conclusive.

In (J3), we address a class of heterogeneous multi-vehicle task assignment and routing problem. We propose two distributed algorithms based on gossip communication: the first algorithm is based on a local exact optimization and the second is based on a greedy heuristic. We consider the case where a set of heterogeneous tasks arbitrarily distributed in a plane has to be serviced by a set of robots, each with a given movement speed and task execution speed. Our goal is to minimize the maximum execution time. We propose two distributed and asynchronous algorithms: the first one is based on the iterative optimization of the local task assignment between pairs of vehicles, the second one is based on local task exchange of assigned tasks, one by one, between couples of vehicles. For both algorithms we provide deterministic bounds to their performance. The proposed approaches are distributed algorithms easy to implement in a networked system and have favorable computational complexity with respect to the ratio k/n between the number of tasks and vehicles.

Bibliography

- (J1) F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, 99(9):1482–1504, 2011
- (J2) S. D. Bopardikar, S. L. Smith, and F. Bullo. On dynamic vehicle routing with time constraints. *IEEE Transactions on Robotics*, 30(6):1524–1532, 2014
- (J3) M. Franceschelli, D. Rosa, C. Seatzu, and F. Bullo. Gossip algorithms for heterogeneous multi-vehicle routing problems. *Nonlinear Analysis: Hybrid Systems*, 10:156–174, 2013
- (J4) P. Agharkar, S. D. Bopardikar, and F. Bullo. Vehicle routing algorithms for radially escaping targets. *SIAM Journal on Control and Optimization*, December 2014. Submitted
- (PhD) P. Agharkar. *Path Planning Algorithms for Autonomous Agents*. PhD thesis, Mechanical Engineering Department, University of California at Santa Barbara, June 2015

3 Coverage control and equitable partitioning

In applications such as environmental monitoring or warehouse logistics a team of robots is asked to perform tasks over a large space. The distributed *environment partitioning problem* consists of designing control and communication laws for individual robots such that the team divides a space into regions in order to optimize the quality of service provided. *Coverage control* additionally optimizes the positioning of robots inside of a region. We consider a comprehensive set of scenarios with different communication protocols and spatial constraints.

In recent work (see C1 and J1 below), we consider the problem of optimal coverage with area-constraints in a mobile multi-agent system. For a planar environment with an associated density function, this problem is equivalent to dividing the environment into optimal subregions such that each agent is responsible for the coverage of its own region. For this work, we design a continuous-time distributed policy which allows a team of agents to achieve a convex area-constrained partition of a convex workspace. Our work is related to the classic Lloyd algorithm, and makes use of generalized Voronoi diagrams. We also discuss practical implementation for real mobile networks.

Bibliography

- (J1) R. Patel, P. Frasca, and F. Bullo. Centroidal area-constrained partitioning for robotic networks. *ASME Journal of Dynamic Systems, Measurement, and Control*, 136(3):031024, 2014
- (C1) R. Patel, P. Frasca, and F. Bullo. Centroidal area-constrained partitioning for robotic networks. In *ASME Dynamic Systems and Control Conference*, Stanford, CA, USA, October 2013

4 Knapsack optimization with application to mixed teams

In a recent journal article (J1) below, we study a class of non-convex optimization problems involving sigmoid functions. Examples of sigmoid utility functions include the correctness of human decisions as a function of the decision time, the effectiveness of human-machine communication as a function of the communication rate, human performance in multiple target search as a function of the search time, advertising response as a function of the investment, and the expected profit in bidding as a function of the bidding amount.

We show that sigmoid functions impart a combinatorial element to the optimization variables and make the global optimization computationally hard. We formulate versions of the knapsack problem, the generalized assignment problem and the bin-packing problem with sigmoid utilities. We merge approximation algorithms from discrete optimization with algorithms from continuous optimization to develop approximation algorithms for these NP-hard problems with sigmoid utilities.

Bibliography

- (J1) V. Srivastava and F. Bullo. Knapsack problems with sigmoid utility: Approximation algorithms via hybrid optimization. *European Journal of Operational Research*, 236(2):488–498, 2014
- (PhD) V. Srivastava. *Stochastic Search and Surveillance Strategies for Mixed Human-Robot Teams*. PhD thesis, Mechanical Engineering Department, University of California at Santa Barbara, December 2012

5 Controllability for complex networks

In a recent technical report we studied the problem of controlling complex networks, that is, the joint problem of selecting a set of control nodes and of designing a control input to steer the network to a target state. For this problem (i) we propose a metric to quantify the difficulty of the control problem as a function of the required control energy, (ii) we derive bounds based on the system dynamics (network topology and weights) to characterize the tradeoff between the control energy and the number of control nodes, and (iii) we propose a distributed strategy with performance guarantees for the control of complex networks. In our strategy we select control nodes by relying on network partitioning, and we design the control input by leveraging optimal and distributed control techniques. Our findings show for instance that (i) if the number of control nodes is constant, then the control energy increases exponentially with the number of the network nodes, (ii) if the number of control nodes is a fixed fraction of the network nodes, then certain networks can be controlled with constant energy independently of the network dimension, and (iii) clustered networks may be easier to control because, for sufficiently many control nodes, the control energy depends only on the controllability properties of the clusters and on their coupling strength. We validate our results with examples from power networks, social networks, and epidemics spreading.

Bibliography

- (J1) F. Pasqualetti, S. Zampieri, and F. Bullo. Controllability metrics, limitations and algorithms for complex networks. *IEEE Transactions on Control of Network Systems*, 1(1):40–52, 2014
- (C1) F. Pasqualetti, S. Zampieri, and F. Bullo. Controllability metrics, limitations and algorithms for complex networks. In *American Control Conference*, pages 3287–3292, Portland, OR, USA, June 2014

6 Cooperative patrolling for robotic and camera networks

Coordinated teams of autonomous agents have recently been used for many tasks requiring repetitive execution, including the monitoring of oil spills, the detection of forest fires, the track of border changes, and the patrol (surveillance) of an environment. The surveillance of an area of interest requires the agents to continuously and repeatedly sweep the environment, and the challenging problem consists in scheduling the agents trajectories so as to optimize a certain performance criteria.

Our research in this direction has considered the patrolling of an area of interest by means of a team of mobile robots and by means of a network of autonomous cameras. For both cases, we define suitable performance functions, we characterize optimal trajectories, and we design centralized and distributed algorithms to steer the agents along optimal trajectories.

Regarding the patrolling problem with mobile robots, in article (J1) below, we assume the robots to be identical and capable of sensing and communicating within a certain spatial range, and of moving with bounded speed. We represent the environment as a graph, in which the vertices correspond to physical and strategically important locations, and in which the edges denote the possibility of moving and communicating between locations. Regarding the performance criteria of a patrolling trajectory, we consider (i) the time gap between any two visits of the same region, called refresh time, and (ii) the time needed to inform the team of robots about an event occurred in the environment, called latency. Our contributions to the patrolling problem with mobile robots are as follows. First, we mathematically formalize our patrolling problem with refresh time and latency cost functions. We show that, in general, designing team trajectories with minimum refresh time is an NP-hard optimization problem, and we identify cases (for instance when the roadmap representing the environment has a chain or tree structure) in which an optimal trajectory can be computed efficiently. Second, we propose different approximate trajectories for general roadmaps, and we characterize the performance of our approximate patrolling trajectories. Third, we develop distributed algorithms for the robots to synchronize along the proposed approximate trajectories. Our algorithms leverage on tools from convex and combinatorial optimization, and rely upon mild robots communication assumptions. Fourth and finally, we validate our findings through simulations and experiments: we use the Player/Stage simulation software to show the effectiveness and the robustness of our second patrolling procedure in a campus environment, and we conduct an experiment with real hardware in an indoor environment with obstacles. The experiments confirm the robustness of the proposed strategies against noise and unmodeled dynamics.

Regarding the patrolling problem with a network of autonomous cameras, in article (J2) below, we consider Pan-Tilt-Zoom (PTZ) cameras installed at important locations. We assume the cameras to move their field of view (f.o.v.) to cooperatively surveil the whole environment. As for the case of robots, we model the environment with a robotic roadmap. We develop algorithms for the cameras to self-organize and to detect intruders in the environment, that appear at arbitrary locations and times. We consider static intruders, which remain stationary, and dynamic intruders, which move to avoid detection, if possible. As performance criteria we consider the worst-case and the average detection times, that is the longest time and the average time needed for the cameras to detect intruders. Our contributions in (J2) are as follows. First, we define the camera surveillance problem for the detection of static and dynamic intruders, and we formalize the worst-case and average detection times of static and dynamic intruders. Second, we exhaustively discuss the case of static intruders. We show that, for tree and ring roadmaps, cameras trajectories with minimum worst-case detection time can be designed by solving a continuous graph partitioning problem. For general cyclic roadmaps, our trajectories based on continuous partitions are proved to be optimal up to a factor 2. Third, for the case of dynamic intruders, we derive a necessary and sufficient condition on the cameras locations for the existence of a trajectory with finite detection time. We focus on ring and tree roadmaps. In particular, for the case of ring roadmaps we design a trajectory with detection time within a factor $3/2$ of optimal. Instead, for tree roadmaps, the performance of our trajectory is within a factor 2 of optimal. Fourth and finally, we consider three different communication models, and we propose distributed algorithms for the cameras for continuous graph partitioning in all these scenarios. In particular: our first algorithm assumes a synchronous mode of operation of the cameras; our second algorithm assumes an asymmetric broadcast communication model and extends the class of block-coordinate descent algorithms to the constrained case; and our third algorithm only requires gossip communication. We prove convergence of all these algorithms, and we analyze their performance in a simulation study.

Bibliography

- (J1) F. Pasqualetti, J. W. Durham, and F. Bullo. Cooperative patrolling via weighted tours: Performance analysis and distributed algorithms. *IEEE Transactions on Robotics*, 28(5):1181–1188, 2012
- (J2) D. Borra, F. Pasqualetti, and F. Bullo. Continuous graph partitioning for camera network surveillance. *Automatica*, 52:227–231, 2014
- (J3) F. Pasqualetti, F. Zanella, J. R. Peters, M. Spindler, R. Carli, and F. Bullo. Camera network coordination for intruder detection. *IEEE Transactions on Control Systems Technology*, 22(5):1669–1683, 2014
- (J4) F. Pasqualetti, D. Borra, and F. Bullo. Consensus networks over finite fields. *Automatica*, 50(2):349–358, 2014
- (J5) S. Susca, P. Agharkar, S. Martínez, and F. Bullo. Synchronization of beads on a ring by feedback control. *SIAM Journal on Control and Optimization*, 52(2):914–938, 2014

- (C1) M. Spindler, F. Pasqualetti, and F. Bullo. Distributed multi-camera synchronization for smart-intruder detection. In *American Control Conference*, pages 5120–5125, Montréal, Canada, June 2012
- (C2) D. Borra, F. Pasqualetti, and F. Bullo. Continuous graph partitioning for camera network surveillance. In *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 228–233, Santa Barbara, CA, USA, September 2012
- (C3) F. Zanella, F. Pasqualetti, R. Carli, and F. Bullo. Simultaneous boundary partitioning and cameras synchronization for optimal video surveillance. In *IFAC Workshop on Distributed Estimation and Control in Networked Systems*, pages 1–6, Santa Barbara, CA, USA, September 2012

7 Distributed optimization and task assignment

The increasing interest in performing complex tasks via multi-agent systems (e.g., sensor and robotic networks) has raised the interest in solving distributed optimization problems. The fundamental paradigms in distributed computation are that: (i) information, relevant for the solution of the problem, is distributed all over a network of processors with limited memory and computation capability, and (ii) the overall computation relies only on local computation and information exchange amongst neighboring processors.

Our research in this direction has focused on two complementary lines. In the first line of research (see journal article (J1) below), we consider a distributed version of linear programs. Optimizing linear objectives over linear constraints takes a central role in the optimization literature and thus deserves particular attention also in distributed computation. Each processor in the network is assigned only the information relative to a subset of the decision variables. The objective is to reach an agreement on a global minimum of the problem, if one exists, or to agree that the problem is either unbounded or infeasible. We show in this paper that the Constraints Consensus algorithm is for certain problems dual to our new algorithm. For the new algorithm, we consider linear programs in equality form, which are allowed to be fully degenerate. The Constraints Consensus can in general not be applied to such problems. With respect to this we propose distributed assignment problems as an important problem class, which can be efficiently solved by our new algorithm. The new formulation allows us to adopt all the tools known in the classical simplex literature to this problem setup. In particular simple and efficient decision rules for the pivot iteration can be provided which are a simple mean to deal with degenerate linear programs. Finally, our new algorithm can deal with infeasible and unbounded problems.

In the second line of research, see journal article (J2) below, we consider distributed approaches for the solution of random convex programs (RCP). RCPs are convex optimization problems with a (usually large) number N of randomly extracted constraints; they arise in several application areas, especially in the context of decision under uncertainty. We here consider a setup in which instances of the random constraints (the scenario) are not held by a single centralized processing unit, but are instead distributed among different nodes of a network. Each node "sees" only a small subset of the constraints, and may communicate with neighbors. The objective is to make all nodes converge to the same solution as the centralized RCP problem. To this end, we develop two distributed algorithms that are variants of the constraints consensus algorithm: the active constraints consensus (ACC) algorithm, and the vertex constraints consensus (VCC) algorithm.

The major contributions of this work are threefold. First, we develop two novel algorithms, ACC and VCC, for the computation of distributed solution to RCP. We show that the ACC algorithm converges to the global solution in finite time, and that it requires almost surely bounded communication at each iteration. We give some numerical evidence of the fact that the ACC algorithm converges in a number of iterations that is linear in the communication graph diameter. We also provide numerical evidence that parallel implementation of the ACC algorithm significantly reduces the computation time over the centralized computation time. As a side result, we show that the ACC algorithm may distributively compute the solution of any convex program, and that it is particularly effective when the dimension of decision variable is small compared with the number of constraints. Second, for the special case when the constraints of the RCP are convex in the uncertain parameters, we develop the vertex constraints consensus (VCC) algorithm. We prove that the VCC algorithm converges to the global solution in a number of iterations equal to the diameter of the communication graph. Moreover, we devise a quantized vertex constraints consensus (qVCC) algorithm in which each node has a bounded communication bandwidth. We provide theoretical bounds on the number of the iterations required for qVCC algorithm to converge. Last, we use the proposed algorithms for distributed classification, distributed estimation, and parallel model predictive control.

Bibliography

- (J1) M. Bürger, G. Notarstefano, F. Bullo, and F. Allgöwer. A distributed simplex algorithm for degenerate linear programs and multi-agent assignment. *Automatica*, 48(9):2298–2304, 2012
- (J2) L. Carlone, V. Srivastava, F. Bullo, and G. C. Calafiore. Distributed random convex programming via constraints consensus. *SIAM Journal on Control and Optimization*, 52(1):629–662, 2014
- (C1) L. Carlone, V. Srivastava, F. Bullo, and G. C. Calafiore. A distributed algorithm for random convex programming. In *Int. Conf. on Network Games, Control and Optimization (NetGCooP)*, pages 1–7, Paris, France, October 2011